

Tutorial for Bayesian forensic likelihood ratio

Niko Brümmer
AGNITIO Research, South Africa

October 2011

1 Introduction

In the Bayesian paradigm for presenting forensic evidence to court, it is recommended that the weight of the evidence be summarized as a *likelihood ratio* (LR) between two opposing hypotheses of how the evidence could have been produced. Such LRs are necessarily based on probabilistic models, the parameters of which may be uncertain. It has been suggested by some authors that the value of the LR, being a function of the model parameters should therefore also be considered uncertain and that this uncertainty should be communicated to the court.

In this tutorial, we consider a simple example of a *fully Bayesian* solution, where model uncertainty is integrated out to produce a value for the LR which is *not* uncertain. We show that this solution agrees with common sense. In particular, the LR magnitude is a function of the amount of data that is available to estimate the model parameters.

Bayesian methods are often criticised because of the difficulty of choosing appropriate priors, especially when the priors are non-informative. We do not deny these difficulties, but the problem is not solved by adopting frequentist methods that effectively sweep the prior under the carpet and pretend it does not exist. In this tutorial we do need to choose a non-informative prior and we choose it by examining the effect it has on the end-result.

We shall reference the following books: E.T. Jaynes, *Probability Theory: The Logic of Science*, Cambridge University Press 2003, which we shall abbreviate as PTLOS; and D.J. Balding, *Weight-of-evidence for Forensic DNA Profiles*, Wiley 2005, abbreviated as WEFDNA.

2 Simplified DNA model

In this tutorial we shall derive the details of how to compute the LR with a *simplified* DNA-like model. The idea is not to provide a recipe that can be used in real forensic DNA analysis, but rather to choose a model that facilitates better understanding of the basic look and feel of a fully Bayesian solution. We need the model to be very simple so that we can perform the Bayesian integrals in closed form. More realistic models would require more complex methods, which would obscure the primary purpose of this tutorial.

We suppose that the *DNA profile* of every individual has K different binary *loci* the *state* of each of which can be either 1 or 0. Every individual is therefore categorized by K binary variables, which gives a total number of 2^K states.¹ We represent a DNA profile by a vector of the form $\mathbf{a} = (a_1, a_2, \dots, a_K)$, where $a_k \in \{0, 1\}$ represents the state of locus k .

We assume that given a DNA sample (either recovered at the crime scene where it was left by the *perpetrator*, or obtained from the *suspect*), the state of each locus may be determined without error.

The main complication is when all suspect and perpetrator loci match, that there is a non-zero probability that some person other than the suspect could have the same DNA profile. To compute this probability, we need to model profile distributions.

3 Profile distribution model

Here we define a generative model that is probably about as simple as it can be. Again, our goal is just to illustrate the basic principles of a fully Bayesian approach to this kind of problem. The goal of this exercise is not to reproduce a realistic DNA model—in real population genetics, the models are more complex.

Let the probability that locus k of a randomly chosen person has state 1 be q_k , and the probability that it has state 0 be $1 - q_k$. According to this model we assume the following independencies:

- The locus states are independent: knowing the state of locus k for one or more individuals, tells us nothing about the states of other loci k' .

¹In real DNA profiling, there are different locus types, with more complex state spaces. For example, STR loci consist of two parts with independent states, one inherited from the father and the other from the mother. Each part has 2 or more states, called *alleles*. DNA profiling technology can detect the state of each part, but does not show which comes from the mother and which from the father.

- For each locus k , the binary state for each person is sampled as an *iid* Bernoulli trial with parameter q_k .

We can collect the locus probabilities in the vector² $\mathbf{q} = (q_1, q_2, \dots, q_K)$. We refer to \mathbf{q} as the *model parameter*, which encodes everything there is to know (under the above modelling assumptions) about how locus states are distributed in the population. The model can be summarized by:

$$P(\mathbf{a}|\mathbf{q}) = \prod_{k=1}^K q_k^{a_k} (1 - q_k)^{1-a_k} \quad (1)$$

which is the probability that a randomly chosen individual has DNA profile $\mathbf{a} = (a_1, a_2, \dots, a_K)$ in a population characterized by the model parameter $\mathbf{q} = (q_1, q_2, \dots, q_K)$. The complication is that we are *not* given \mathbf{q} . Its value has to be inferred from prior assumptions and from data.

4 Inferring the model parameter

We do a Bayesian inference for the value of \mathbf{q} , by computing a posterior distribution.

4.1 Prior

As prior for q_k , we assign a *beta distribution*. This choice has a threefold motivation: (i) The beta distribution is a conjugate prior for this problem, which allows for closed-form Bayesian calculations. (ii) It is commonly used in forensic DNA practice. (iii) It is general enough to include various non-informative priors, which will be of special interest to us.

We assign independently for each q_k a beta distribution with hyper-parameter $\pi_k = (\alpha_k, \beta_k)$, so that:

$$P(\mathbf{q}|\boldsymbol{\pi}) = \prod_{k=1}^K \text{Beta}(q_k|\alpha_k, \beta_k) \quad (2)$$

$$= \prod_{k=1}^K \frac{q_k^{\alpha_k-1} (1 - q_k)^{\beta_k-1}}{B(\alpha_k, \beta_k)} \quad (3)$$

²Note that the elements of \mathbf{q} usually do not sum to one. These are K independent probabilities, not one K -ary categorical distribution.

where we have defined $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots, \pi_K)$. The normalization constant of the beta distribution is given by the *beta function*, defined as:

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)} = \int_0^1 q^{\alpha-1}(1-q)^{\beta-1} dq \quad (4)$$

where Γ is the gamma function.

For the beta distribution to be normalized, we need $\alpha_k, \beta_k > 0$ and unless stated otherwise, we shall assume this condition holds for all our calculations below. In places, we will however consider the limit as $\alpha_k = \beta_k \rightarrow 0$. When we do this, we will follow the advice of PTLOS and complete the whole calculation under the assumption $\alpha_k, \beta_k > 0$ and apply the limit only to the final result.

4.1.1 Non-informative priors

If we want to use a non-informative prior, we let $\alpha = \alpha_k = \beta_k$ by symmetry, and we can choose some α , for example in the range $0 < \alpha \leq 1$. The case $\alpha \rightarrow 0$ is called the *Haldane* prior, the case $\alpha = 0.5$ is the *Jeffreys* prior and $\alpha = 1$ is the *Laplace* prior.

The Haldane prior is *flat* in the sense that the probability density for $\log \frac{q}{1-q}$ is uniform, but since this reparametrization of q covers the whole real line, this prior is improper.

The Jeffreys prior is flat in the sense that the probability density for $\arcsin(2q - 1)$ is uniform between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$.

The Laplace prior is flat in the sense that the probability density for q is uniform between 0 and 1.

As these names show, different workers in probability theory have arrived at different conclusions about which prior should be used to encode non-informativeness about the Bernoulli model parameter. To make our calculations concrete, we will have to make a definite choice of prior. We shall solve this problem in a later section, by examining the effect of the prior on the end-result of our calculation.

4.1.2 Informative prior

In forensic DNA³ it is customary to reparametrize the beta prior as:

$$\alpha_k = \frac{1 - \theta}{\theta} p_k, \quad \beta_k = \frac{1 - \theta}{\theta} (1 - p_k) \quad (5)$$

³See WEFDNA pp. 63-64.

where $0 < p_k < 1$ and $0 < \theta < 1$. Here θ is known as the *population structure parameter*. With this parametrization, $\text{Beta}(q_k|\alpha_k, \beta_k)$ has the following mean and variance:

$$\langle q_k \rangle = \frac{\alpha_k}{\alpha_k + \beta_k} = p_k \quad \langle (q_k - p_k)^2 \rangle = \theta p_k (1 - p_k) \quad (6)$$

For small values of θ , one obtains an *informative* prior, with a small variance and a sharp peak near p_k . In the extreme as $\theta \rightarrow 0$, we get a strongly informative prior, which will override contributions made by finite data and therefore asserts $q_k = p_k$.

For the case $p_k = \frac{1}{2}$ and $\theta = \frac{1}{2\alpha+1} \geq \frac{1}{3}$, we recover the above-mentioned non-informative priors: Laplace at $\theta = \frac{1}{3}$, Jeffreys at $\theta = \frac{1}{2}$ and in the extreme as $\theta \rightarrow 1$, the Haldane prior, which gives maximum weight to the data. These effects will be shown below.

4.2 Database

We make provision in our calculation to optionally use a database of examples to help us infer values for \mathbf{q} . Let $\mathbf{A} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_L)$ be a database of DNA profiles for L different individuals, where the profile for individual ℓ is $\mathbf{a}_\ell = (a_{1\ell}, a_{2\ell}, \dots, a_{K\ell})$ and where $a_{k\ell} \in \{0, 1\}$ is the binary state of locus k of individual ℓ . We assume the DNA profiles in \mathbf{A} :

- have been sampled iid from the same population as the suspect and perpetrator and are therefore relevant to inferring the parameter \mathbf{q} ,
- but the individuals are distinct from the suspect and the perpetrator.

Our calculations will allow for the case of the empty database, where $L = 0$.

4.3 Likelihood

Because of our independence assumptions in the model, the likelihood for \mathbf{q} , given the database \mathbf{A} is:

$$P(\mathbf{A}|\mathbf{q}) = \prod_{\ell=1}^L \prod_{k=1}^K q_k^{a_{k\ell}} (1 - q_k)^{1-a_{k\ell}} \quad (7)$$

$$= \prod_{k=1}^K q_k^{n_k} (1 - q_k)^{L-n_k} \quad (8)$$

where $n_k = \sum_{\ell=1}^L a_{k\ell}$ is the number of times locus k has state 1 and $L - n_k$ is the number of times it has state 0.

4.4 Posterior

We can now infer the value of \mathbf{q} by computing the posterior:

$$P(\mathbf{q}|\mathbf{A}, \boldsymbol{\pi}) = \frac{P(\mathbf{q}|\boldsymbol{\pi})P(\mathbf{A}|\mathbf{q})}{\int P(\mathbf{q}'|\boldsymbol{\pi})P(\mathbf{A}|\mathbf{q}') d\mathbf{q}'} \quad (9)$$

$$= \prod_{k=1}^K \frac{q_k^{\alpha_k+n_k-1}(1-q_k)^{\beta_k+L-n_k-1}}{\int_0^1 q_k'^{\alpha_k+n_k-1}(1-q_k')^{\beta_k+L-n_k-1} dq_k'} \quad (10)$$

$$= \prod_{k=1}^K \text{Beta}(q_k|\alpha_k+n_k, \beta_k+L-n_k) \quad (11)$$

where the integral in the denominator was solved by inspection, by recognizing the numerator as another beta distribution. This is due to the fact that the beta distribution is conjugate to the Bernoulli likelihood and therefore should result in a beta posterior. Notice that if the database is empty, then $n_k = L = 0$ and the posterior is just the prior.

The prior parameters α_k and β_k play the same roles mathematically as the event counts n_k and $L - n_k$ and are consequently referred to as *pseudo-counts*. The total pseudo count, $\alpha + \beta$ can be interpreted as the *size* of some pseudo database, which is then effectively pooled with \mathbf{A} by the additions in (11).

In the alternative prior parametrization, $\frac{1-\theta}{\theta}p_k$ and $\frac{1-\theta}{\theta}(1-p_k)$ are the pseudo counts and $\frac{1-\theta}{\theta}$ is the size of the pseudo database.

The posterior $P(\mathbf{q}|\mathbf{A}, \boldsymbol{\pi})$ represents our total state of knowledge about \mathbf{q} and can be used in all calculations in place of the unknown \mathbf{q} .

5 Forensic LR

We are given two DNA profiles: One for the *suspect*, $\mathbf{s} = (s_1, s_2, \dots, s_K)$ and one for the *perpetrator*, $\mathbf{r} = (r_1, r_2, \dots, r_K)$. We work with two hypotheses and assume they are the *only* possible explanations for the observed data \mathbf{s}, \mathbf{r} :

- The *prosecution hypothesis*, H_p , asserts that suspect and perpetrator are the same person.
- The *defence hypothesis* H_d , asserts that they are different individuals.

Below we compute the likelihoods under each hypothesis. For now, we assume that if they don't match, $\mathbf{r} \neq \mathbf{s}$, then in the absence of DNA measurement errors, this proves deductively that H_d is true and H_p is false.

In the matched case, $\mathbf{r} = \mathbf{s}$, however, we need probabilistic reasoning. The most natural way to do this would be to compute the *posterior*,

$$P(H_p|\mathbf{r}, \mathbf{s}, \mathbf{A}, \boldsymbol{\pi}, \Pi) = 1 - P(H_d|\mathbf{r}, \mathbf{s}, \mathbf{A}, \boldsymbol{\pi}, \Pi) \quad (12)$$

where we have introduced the *prior for the prosecution hypothesis*,

$$\Pi = P(H_p|\Pi) = 1 - P(H_d|\Pi) \quad (13)$$

which is assigned by a reasoning process not involving DNA profiles. However, in the Bayesian paradigm for presenting evidence in court one equivalently considers the *posterior odds for H_p against H_d* , which can be separated⁴ into two factors: *likelihood ratio* and *prior odds*, respectively representing the contributions of the DNA analysis and all other evidence not related to DNA:

$$\frac{P(H_p|\mathbf{r}, \mathbf{s}, \mathbf{A}, \boldsymbol{\pi}, \Pi)}{P(H_d|\mathbf{r}, \mathbf{s}, \mathbf{A}, \boldsymbol{\pi}, \Pi)} = \text{LR} \frac{\Pi}{1 - \Pi} \quad (14)$$

where

$$\text{LR} = \frac{P(\mathbf{r}, \mathbf{s}|H_p, \mathbf{A}, \boldsymbol{\pi})}{P(\mathbf{r}, \mathbf{s}|H_d, \mathbf{A}, \boldsymbol{\pi})} \quad (15)$$

is referred to as *the likelihood ratio*. It is then recommended that the end-goal of the forensic DNA analysis is to compute LR, which can be done independently of Π . We derive expressions for both likelihoods below and then form the ratio.

Finally, notice that if $\text{LR} = 1$, then the DNA analysis is *completely non-informative* about H_p versus H_d : in this case the posterior (odds) is the same as the prior (odds).

5.1 Prosecution likelihood

Under the prosecution hypothesis, \mathbf{r} and \mathbf{s} come from the same individual, so that $P(\mathbf{r}, \mathbf{s}|H_p, \mathbf{q}) = \delta(\mathbf{r}, \mathbf{s})P(\mathbf{s}|\mathbf{q})$, where $\delta(\mathbf{r}, \mathbf{r}) = 1$, or $\delta(\mathbf{r}, \mathbf{s}) = 0$ if $\mathbf{r} \neq \mathbf{s}$. Since we are not given \mathbf{q} , but instead we are given the prior $\boldsymbol{\pi}$ and the database \mathbf{A} , we must condition on what we have and instead compute:

$$P(\mathbf{r}, \mathbf{s}|H_p, \boldsymbol{\pi}, \mathbf{A}) = \delta(\mathbf{r}, \mathbf{s})P(\mathbf{s}|\boldsymbol{\pi}, \mathbf{A}) \quad (16)$$

⁴In the real world, this simple factorization applies only in a limited number of cases. If different alternative culprits, with different levels of relatedness to the suspect are considered, somewhat more general formulas have to be used, as explained in WEFDNA.

where

$$P(\mathbf{s}|\boldsymbol{\pi}, \mathbf{A}) = \int_0^1 \int_0^1 \cdots \int_0^1 P(\mathbf{s}|\mathbf{q})P(\mathbf{q}|\boldsymbol{\pi}, \mathbf{A}) dq_1 dq_2 \cdots dq_K \quad (17)$$

$$= \int_Q P(\mathbf{s}|\mathbf{q})P(\mathbf{q}|\boldsymbol{\pi}, \mathbf{A}) d\mathbf{q} \quad (18)$$

where Q is short-hand for the K -cube over which we are integrating. Note $P(\mathbf{s}|\boldsymbol{\pi}, \mathbf{A})$ is called the *predictive distribution* for \mathbf{s} , because it predicts the value of an as yet unseen profile, given that we have already seen the profiles in \mathbf{A} . Again by virtue of the conjugate prior, the predictive distribution can be found in closed form:

$$P(\mathbf{s}|\boldsymbol{\pi}, \mathbf{A}) = \int_Q P(\mathbf{s}|\mathbf{q})P(\mathbf{q}|\boldsymbol{\pi}, \mathbf{A}) d\mathbf{q} \quad (19)$$

$$= \prod_{k=1}^K \int_0^1 q_k^{s_k} (1 - q_k)^{1-s_k} \frac{q_k^{\alpha_k+n_k-1} (1 - q_k)^{\beta_k+L-n_k-1}}{B(\alpha_k + n_k, \beta_k + L - n_k)} dq_k \quad (20)$$

$$= \prod_{k=1}^K \frac{\int_0^1 q_k^{\alpha_k+s_k+n_k-1} (1 - q_k)^{\beta_k+L+1-s_k-n_k-1} dq_k}{B(\alpha_k + n_k, \beta_k + L - n_k)} \quad (21)$$

$$= \prod_{k=1}^K \frac{B(\alpha_k + s_k + n_k, \beta_k + L + 1 - s_k - n_k)}{B(\alpha_k + n_k, \beta_k + L - n_k)} \quad (22)$$

$$= \prod_{k=1}^K P(s_k|\pi_k, n_k, L) \quad (23)$$

Now we can expand the beta functions in terms of gamma functions and simplify the ratios of gammas with the identity $\Gamma(x+1) = x\Gamma(x)$, to find the predictive probability:⁵

$$P(s_k = 1|\pi_k, n_k, L) = \frac{\alpha_k + n_k}{\alpha_k + \beta_k + L} \quad (24)$$

$$= \frac{(1 - \theta)p_k + \theta n_k}{(1 - \theta) + \theta L} \quad (25)$$

For the informative prior case, notice that θ gives *interpolation* weights between data and the prior parameter p_k . At the one extreme if $\theta \rightarrow 1$ (Haldane prior), we disregard the prior parameter p_k and end up with just the data proportion $\frac{n_k}{L}$. At the other extreme if $\theta = 0$, we disregard the data \mathbf{A} and end up with the prior parameter p_k . (If we use the non-informative Laplace

⁵Notice (25) agrees with equation 5.6 on page 64 in WEFDNA.

prior, with $\alpha_k = \beta_k = 1$, then (24) is known as Laplace's *rule of succession*.) Finally, the predictive probability⁶ for the event $s_k = 0$ is:

$$P(s_k = 0 | \pi_k, n_k, L) = \frac{\beta_k + L - n_k}{\alpha_k + \beta_k + L} \quad (26)$$

$$= \frac{(1 - \theta)(1 - p_k) + \theta(L - n_k)}{(1 - \theta) + \theta L} \quad (27)$$

Note that even for an empty database ($L = n_k = 0$), our assumption $\alpha_k, \beta_k > 0$ guarantees non-zero predictive probabilities.

5.2 Defence likelihood

Under the defence hypothesis, \mathbf{r} and \mathbf{s} come from different individuals and their probabilities are independent given \mathbf{q} , so that $P(\mathbf{r}, \mathbf{s} | \mathbf{q}) = P(\mathbf{r} | \mathbf{q})P(\mathbf{s} | \mathbf{q})$. However, \mathbf{q} is not given, so the independence no longer holds: knowledge of one profile changes the probability for \mathbf{q} , which in turn changes the probability for the other profile. This dependency is automatically taken care of by applying the rules of probability theory by integrating out the unknown \mathbf{q} :

$$P(\mathbf{r}, \mathbf{s} | H_d, \boldsymbol{\pi}, \mathbf{A}) = \int_Q P(\mathbf{r} | \mathbf{q})P(\mathbf{s} | \mathbf{q})P(\mathbf{q} | \boldsymbol{\pi}, \mathbf{A}) d\mathbf{q} \quad (28)$$

$$= \prod_{k=1}^K \frac{B(\alpha_k + s_k + r_k + n_k, \beta_k + L + 2 - r_k - s_k - n_k)}{B(\alpha_k + n_k, \beta_k + L - n_k)} \quad (29)$$

$$= \prod_{k=1}^K P(r_k, s_k | \pi_k, n_k, L) \quad (30)$$

where we can expand and simplify again to find the predictive probability:

$$P(r_k = s_k = 1 | \pi_k, n_k, L) = \frac{\alpha_k + n_k}{\alpha_k + \beta_k + L} \frac{\alpha_k + n_k + 1}{\alpha_k + \beta_k + L + 1} \quad (31)$$

$$= P(r_k = 1 | \pi_k, n_k, L)P(s_k = 1 | \pi_k, n_k + 1, L + 1) \quad (32)$$

Notice the similarity between the two factors in the RHS: the right factor is obtained from the left by adding 1's to the observation counts. Notice also that if $\alpha + n_k \gg 1$, then $P(r_k = s_k = 1 | \pi_k, n_k, L) \approx P(r_k = 1 | \pi_k, n_k, L)P(s_k = 1 | \pi_k, n_k, L)$, making the two events *almost* independent.

⁶Notice $P(s_k = 0 | \alpha_k, \beta_k, n_k, L) + P(s_k = 1 | \alpha_k, \beta_k, n_k, L) = 1$.

The probability for the other event of interest⁷ is obtained similarly as:

$$P(r_k = s_k = 0 | \pi_k, n_k, L) = P(r_k = 0 | \pi_k, n_k, L) P(s_k = 0 | \pi_k, n_k, L + 1) \quad (33)$$

5.3 LR

Forming the likelihood-ratio, we find:

$$\text{LR} = \frac{P(\mathbf{r}, \mathbf{s} | H_p, \boldsymbol{\pi}, \mathbf{A})}{P(\mathbf{r}, \mathbf{s} | H_d, \boldsymbol{\pi}, \mathbf{A})} = \prod_{k=1}^K \text{LR}_k(r_k, s_k) \quad (34)$$

where

$$\text{LR}_k(r, s) = \frac{\delta(r, s) P(s | \pi_k, n_k, L)}{P(r, s | \pi_k, n_k, L)} \quad (35)$$

$$= \frac{\delta(r, s) P(s | \pi_k, n_k, L)}{P(s, s | \pi_k, n_k, L)} \quad (36)$$

$$= \frac{\delta(r, s) P(s | \pi_k, n_k, L)}{P(s | \pi_k, n_k, L) P(s | \pi_k, n_k + s, L + 1)} \quad (37)$$

$$= \frac{\delta(r, s)}{P(s | \pi_k, n_k + s, L + 1)} \quad (38)$$

More explicitly, for the mismatched cases we have

$$\text{LR}_k(0, 1) = \text{LR}_k(1, 0) = 0 \quad (39)$$

and for the matched cases we have

$$\text{LR}_k(1, 1) = \frac{\alpha_k + \beta_k + L + 1}{\alpha_k + n_k + 1}, \quad \text{LR}_k(0, 0) = \frac{\alpha_k + \beta_k + L + 1}{\beta_k + L + 1 - n_k} \quad (40)$$

or, with the other prior parametrization:

$$\text{LR}_k(1, 1) = \frac{(1 - \theta) + \theta(L + 1)}{(1 - \theta)p_k + \theta(n_k + 1)} \quad (41)$$

and

$$\text{LR}_k(0, 0) = \frac{(1 - \theta) + \theta(L + 1)}{(1 - \theta)(1 - p_k) + \theta(L + 1 - n_k)} \quad (42)$$

Notice again, that θ interpolates between data and the prior parameter p_k . The minimum value (for the matched case $r_k = s_k$) is 1. This is a consequence of the error-free measurement assumption. If non-zero error probabilities were considered, values of less than 1 would be possible.

⁷We don't need the events (0, 1) and (1, 0) here, because we are interested in the case where profiles match.

6 Plug-in recipe

In this section, we shall refer to:

- One or more *reference populations*, from which one or more databases are drawn to help to estimate the parameters p_k and θ for an *informative* prior.
- The *relevant population*, from which the suspect and perpetrator were drawn.

In the general case, all these populations are assumed different from each other in the sense that locus state *frequencies may differ* between them.

For forensic DNA applications, WEFDNA motivates a *plug-in* recipe to compute the LR, where values for θ and the p_k are point-estimates made from one or more reference databases. In this recipe, the p_k are representative of the frequencies in the *reference* populations, while the value of θ is chosen to reflect by how much the corresponding frequencies in the *relevant* population may differ. Small values of θ encode small expected differences and larger values encode larger expected differences. WEFDNA motivates for values in the range $1\% \leq \theta \leq 5\%$ to be used for most applications.

Our database \mathbf{A} , as defined in section 4.2, is assumed to be drawn from the *relevant* population, but in the usual forensic scenario, additional profiles from the relevant population are not available. In our notation, this means \mathbf{A} is empty.

In summary, in the WEFDNA plug-in recipe we set $L = n_k = 0$, the p_k are generally different from $\frac{1}{2}$ and θ is smallish. This forms an *informative* prior for the q_k . This gives, for $r_k = s_k$:

$$\text{LR}_k(1, 1) = \frac{1}{(1 - \theta)p_k + \theta} \quad \text{LR}_k(0, 0) = \frac{1}{(1 - \theta)(1 - p_k) + \theta} \quad (43)$$

7 Fully Bayesian recipe

Now we turn to the main purpose of this document, namely to explore a fully Bayesian recipe, where we start with a *non-informative* prior and use only the given data, $\mathbf{A}, \mathbf{r}, \mathbf{s}$, to infer the model parameter.

It must be emphasized that this fully Bayesian recipe cannot be used as is to replace the plug-in recipe, because here we use the luxury of database \mathbf{A} , sampled from the relevant population. As noted above, in a realistic scenario, we do not have this luxury: instead we have to make do with data sampled from some other, somewhat different, reference population. Although a fully

Bayesian recipe could in principle be derived for this more realistic scenario, this would come at the cost of a considerable increase in both conceptual difficulties as well as computational complexity.

In this section, therefore we assume we *do* have a database, \mathbf{A} , sampled from the relevant database and the only difficulty that remains is to choose the non-informative prior.

7.1 Which prior?

We are now faced with making a choice amongst the different flavours of non-informative priors. That is, we have to choose α_k and β_k , or equivalently p_k and θ .

We concede that we are choosing a prior under the perhaps arbitrary constraint that it should be a beta distribution. A more thorough motivation for the prior should perhaps involve solving functional equations in the style of PTLOS. We feel however that the beta distribution already provides a rich enough space for the choice of prior. Moreover, as mentioned above, the non-informative Haldane, Jeffreys and Laplace priors all members of the beta family.

To start, we motivate the choice $\alpha = \alpha_k = \beta_k$, or equivalently $p_k = \frac{\alpha_k}{\alpha_k + \beta_k} = \frac{1}{2}$. Before we have seen any data, all loci are on an equal footing, so that the priors for all k must be the same. Next consider a database \mathbf{A} with an equal number of 0's and 1's for some locus k , so that $n_k = L - n_k$. In this situation, there is no reason to prefer one state to the other, so that the model parameter posterior should satisfy the symmetry condition: $P(q_k | \alpha_k, \beta_k, L, n_k) = P(1 - q_k | \alpha_k, \beta_k, L, n_k)$, which is obtained at $\alpha_k = \beta_k$. Another way to see this is simply to require $\text{LR}_k(0, 0) = \text{LR}_k(1, 1)$ when $n_k = L - n_k$.

Now we have $p_k = \frac{1}{2}$ and we still need to choose θ . To do this, consider the case of the *empty* database, with $L = n_k = 0$, for which case we still want our recipe to give a sensible answer. Now (41) and (42) give:

$$\text{LR}_k(1, 1) = \text{LR}_k(0, 0) = \frac{1}{(1 - \theta)^{\frac{1}{2}} + \theta} = \frac{2}{1 + \theta} \quad (44)$$

When \mathbf{A} is empty, we now argue that we don't even know whether the locus state varies in the population. So we are not justified in concluding that the match at the locus modifies the probabilities for H_p vs H_d . If we maximize θ at the limit $\theta \rightarrow 1$, then we obtain the non-informative value of $\text{LR}_k = 1$, so that the DNA evidence is effectively *disregarded*.

7.2 Analysis

Here we analyse the behaviour of $\text{LR}_k(r_k, s_k)$, when $r_k = s_k$ and $\theta = 1$. We get:

$$\text{LR}_k(1, 1) = \frac{L + 1}{n_k + 1}, \quad \text{LR}_k(0, 0) = \frac{L + 1}{L + 1 - n_k} \quad (45)$$

We make several observations:

- The matched likelihood ratios are bounded: $1 \leq \text{LR}_k(s, s) \leq L + 1$. We have already commented on the lower bound. The upper bound is determined by the database size, L . This makes intuitive sense, the larger the database, the more our maximum confidence grows. Note however, that this maximum should be a relatively rare occurrence, as shown below.
- For an empty database, if $L = n_k = 0$, then as discussed, $\text{LR}_k(1, 1) = \text{LR}_k(0, 0) = 1$.
- For a non-empty database, as long as a locus k has the same state in all of the observed data, $\mathbf{A}, \mathbf{r}, \mathbf{s}$, then the LR is *still* unity: If $n_k = L$, then $\text{LR}_k(1, 1) = 1$ and if $n_k = 0$, then $\text{LR}_k(0, 0) = 1$.
- Conversely, for a given database size L , the maximum LR value is reached when the locus state observed in $s_k = r_k$ has never been observed in \mathbf{A} . This implies the trait shared by the suspect and perpetrator is *rare*. The larger the database size, L , the more we are convinced of the rarity and the more we are convinced of the identity of suspect and perpetrator.
- For a large database, where both $n_k \gg 1$ and $L - n_k \gg 1$, the likelihood ratio for $s_k = r_k$ is the inverse of the frequency of the corresponding event in the database: $\text{LR}_k(1, 1) \approx \frac{L}{n_k}$ and $\text{LR}_k(0, 0) \approx \frac{L}{L - n_k}$.

We can briefly compare this recipe to a very naive recipe, where we simply assign $q_k = \frac{n_k}{L}$, irrespective of the size of the database. This would give $\text{LR}(1, 1) = \frac{L}{n_k}$ and $\text{LR}_k(0, 0) = \frac{L}{L - n_k}$. This agrees with the last case above of the Bayesian recipe, but in any other cases it could give overconfident results. In particular, if $n_k = 0$, or $n_k = L$, one could get infinite LR values, which would be ridiculous in the extreme if $L = 1$. The fully Bayesian recipe agrees with the naive recipe when data is plentiful, but continues to give sensible answers even when the data gets scarce to the point of vanishing.

7.2.1 Comment on Haldane prior

With a more realistic DNA model, where each STR locus has two independent sides (paternal and maternal), we can gain some extra insight into the nature of the Haldane prior. In this case, it can be shown (WEFDNA, section 6.2.2) that when $L = 0$, the LR for a locus can nevertheless reach a maximum of 3. If the paternal and maternal sides are the same, then we get LR=1, but if they are different, we get LR=3. From this fact and the third bullet above, we learn that:

The LR at locus k becomes non-informative ($LR_k = 1$) under the Haldane prior, if and only if no state change has been observed at locus k in *all* of the data, $\mathbf{A}, \mathbf{r}, \mathbf{s}$.

One may argue that loci used for forensic DNA profiling have been chosen for the purpose of giving good discrimination between individuals, precisely because they *do* vary appreciably between individuals and that therefore the Haldane prior is too extreme. However, we are concerned here with *sub-populations*, about which we cannot assume that every locus is informative—it may well be that a certain locus is constant over the whole sub-population. We therefore argue that the behaviour of the Haldane prior is appropriate: the LR for a locus remains non-informative ($LR_k = 1$), until we have observed at least one state change in our data.